EPOXI Uplink Array Experiment of June 27, 2008

Victor A. Vilnrotter,* Philip C. Tsao,* Dennis K. Lee,* Timothy P. Cornish,† Leslie Paal,† and Vahraz Jamnejad†

Uplink array technology is currently being developed for NASA's Deep Space Network (DSN) to provide greater range and data throughput for future NASA missions, including manned missions to Mars and exploratory missions to the outer planets, the Kuiper Belt, and beyond. The DSN uplink arrays employ N microwave antennas transmitting at 7.2 GHz (X-band) to produce signals that add coherently at the spacecraft, hence providing a power gain of N^2 over a single antenna. This gain can be traded off directly for an N^2 higher data rate at a given distance such as Mars, providing, for example, high-definition video broadcast from Earth to a future human mission, or it can provide a given data rate for commands and software uploads at a distance N times greater than would be possible with a single antenna. The uplink arraying concept has been recently demonstrated using the three operational 34-m antennas of the Apollo Complex at the Goldstone Deep Space Communications Complex in California, which transmitted arrayed signals to the EPOXI spacecraft (an acronym formed from EPOCh and DIXI: Extrasolar Planetary Observation and Characterization and Deep Impact Extended Investigation). Both two-element and three-element uplink arrays were configured, and the theoretical array gains of 6 dB and 9.5 dB, respectively, were demonstrated experimentally. This required initial phasing of the array elements, the generation of accurate frequency predicts to maintain phase from each antenna despite relative velocity components due to Earth rotation and spacecraft trajectory, and monitoring of the ground-system phase for possible drifts caused by thermal effects over the 16-km fiber-optic signal distribution network. This article provides a description of the equipment and techniques used to demonstrate the uplink arraying concept in a relevant operational environment. Data collected from the EPOXI spacecraft are also analyzed to verify array calibration, array gain, and system stability over the entire five-hour duration of this experiment.

I. Introduction

Coherent arraying of antennas transmitting uplink signals to a spacecraft is a novel concept that will greatly increase NASA's deep-space communications capabilities in the future. Typically, deep-space missions require the capability to command the spacecraft right after

* Communication Architectures and Research Section.
† Communications Ground Systems Section.

The research described in this publication was carried out by the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.

launch, during cruise, and after encountering the target, in order to provide two-way communication and ranging. In addition, in-flight reconfiguration via software uploads may be required to accommodate unforeseen changes in mission objectives. The use of antenna arrays enables much greater data rates, greater effective operating distance, and cost-effective scaling for more demanding future missions through a highly flexible design philosophy, via the inherently parallel architecture of antenna arrays.

Currently, uplink command is performed routinely using operational 34- and 70-m antennas, but fundamental limitations in practically attainable antenna aperture, transmitter power, and safety constraints relegate single-antenna uplink to modest distances and command rates. This has implications not only in the operational range of single-aperture uplink antennas, but also for the time it takes to regain communications with a spacecraft in case of emergency, where the location of the spacecraft may not be known exactly. As described in Section II, arrays of a few 34-m (or similar size) antennas can perform a very efficient search over the uncertainty region, enabling rapid recovery of communications with a lost spacecraft.

On June 27, 2008, NASA's Deep Space Network (DSN) configured an array of three 34-m antennas located at the Apollo Station to transmit coherently combined signals to the EPOXI (Extrasolar Planetary Observation and Characterization [EPOCh]/Deep Impact Extended Investigation [DIXI]) spacecraft, yielding nine times greater signal power than previously possible with a single antenna. Prior to this experiment, only single 34-m or 70-m antennas had been used to transmit uplink signals. Successful coherent arraying of three DSN antennas is a major step towards demonstrating future uplink array capabilities, greatly extending the DSN's reach into deep space.

During this experiment, 7.2-GHz (X-band) signals were combined from three 34-m antennas at the Goldstone Deep Space Communications Complex (DSCC) in the Mojave Desert in California, and transmitted to EPOXI. The EPOXI spacecraft received the combined signals, measured the combined power levels, and verified the predicted array power gains of 6 dB and 9.5 dB over a single antenna, for both two- and three-antenna arrays. In addition, non-operational (or No-Op) test messages that resemble real uplink commands (but the commands are ignored by the spacecraft) were transmitted, all of which were received and acknowledged by EPOXI.

Operational uplink array calibration generally takes place before each spacecraft track, after which the array is required to maintain calibration. After uplink array calibration is completed, local measurements of antenna phase are employed to minimize instrumental and fiber-optic phase drift. The uplink arraying functions demonstrated and evaluated during this experiment include:

(1) Compensation for differential Doppler on the transmitted X-band uplink carriers.

(2) "Phase-ramp" calibration of two- and three-antenna arrays.

(3) Spacecraft configuration to deliver automatic gain control (AGC) readings as part of engineering data.

(4) AGC readings obtained in real-time from downlink telemetry and processing of engineering data to extract the sequence of AGC readings.

(5) Real-time analysis of spacecraft data to estimate the optimum uplink array differential phase.

(6) Application of optimum differential phase to the uplink signals.

(7) Real-time monitoring of phase drifts in the ground equipment.

(8) Transmission of realistic No-Op commands to EPOXI (No-Op commands have the same subcarrier and modulation structure as real operational uplink commands, but the spacecraft takes no action).

II. EPOXI Uplink Array Experiment

The goals of this experiment were to illuminate the EPOXI spacecraft simultaneously with two and three X-band carriers, using 34-m beam-waveguide (BWG) antennas at the Goldstone Apollo Complex (DSS-24, DSS-25, and DSS-26), maximize the combined power of the spacecraft, determine the stability of the arrayed carrier power at the spacecraft, and transmit realistic No-Op commands to the spacecraft at the maximum possible rate. The actual physical configuration of the Apollo Complex antennas is shown in Figure 1, where DSS-26 is closest, and DSS-24 farthest along a north–northwest baseline.

Figure 1. The Apollo Complex of 34-m BWG antennas at the Goldstone DSCC.

A conceptual diagram of the experimental setup is shown in Figure 2, similar to the previous experiment carried out earlier with the Mars Global Surveyor spacecraft, and reported in [1]. The central antenna (DSS-25) was designated as the reference antenna during this experiment, and was the only antenna configured to receive EPOXI downlink telemetry and transmit X-band uplink. The telemetry contained 5-s updates of the Small Deep Space Transponder (SDST) wide-band AGC reading (WBAGC), and narrow-band AGC readings (NBAGC), as well as 5-s updates of the coherent lock accumulator (CLA), which provided the most accurate estimate of received carrier power. The other two antennas were employed in transmit mode only.

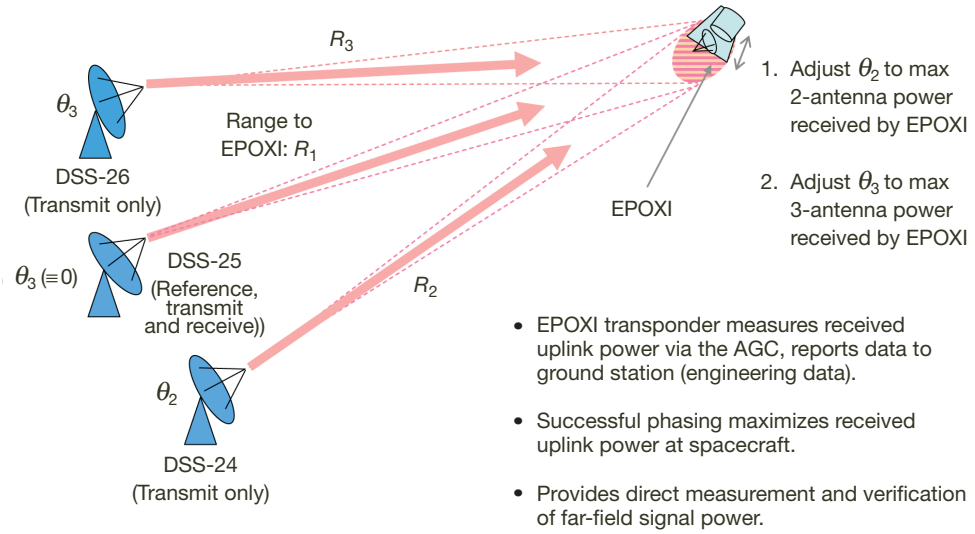

Figure 2. EPOXI Uplink Array Experiment conceptual diagram.

Differential Doppler caused by Earth rotation was removed using extremely accurate frequency predicts developed especially for uplink array applications, and applied to the exciters at Signal Processing Center (SPC)–10 prior to transmission. Examples of the steady-state Doppler-compensated two- and three-antenna power distributions in the vicinity of the spacecraft are shown in the simulated far-field patterns of Figure 3. The actual antenna positions at the Apollo Complex were used to calculate these patterns, for a target located at zenith. These patterns were calculated at the distance of the Moon, and previously used to explain the effects of array illumination on a selected lunar target [2]. In all cases, the arrayed power distribution in the far-field is the product of the primary antenna pattern and the interference pattern generated by two or three point sources located at the antenna phase centers.

It can be seen from Figure 3 that far-field patterns of arrays composed of a few 34-m antennas have an interesting property that could simplify the search for lost spacecraft in future emergency communications applications. For two-antenna arrays, the location of the array peaks and nulls can be interchanged within the entire single-antenna beam, simply by adding 180 degrees to the phase of one of the antennas. Therefore, the entire 34-m beamwidth can be covered with just one additional operation, as discussed further in Section V.

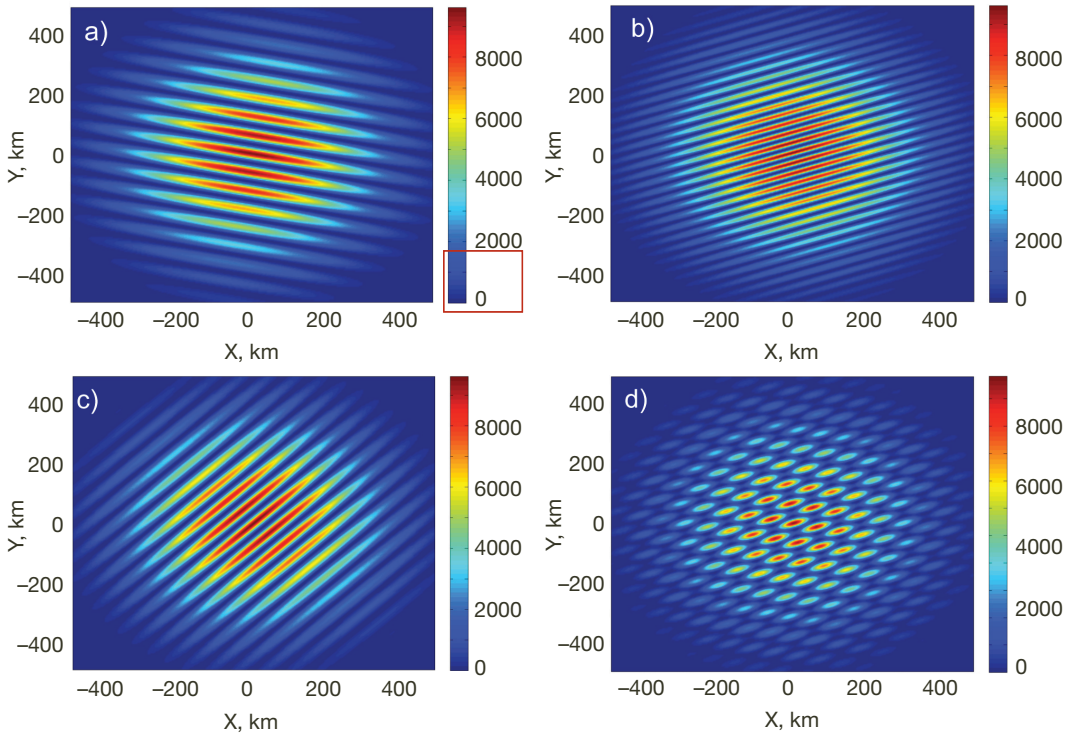

Figure 3. Far-field power distributions generated by two- and three-element arrays of the Apollo Complex, at the distance of the Moon, looking directly overhead. (a) DSS-25/DSS-26 baseline, (b) DSS-24/DSS-26 baseline, (c) DSS-24/DSS-25 baseline, (d) three-antenna far-field intensity pattern. Color code is proportional to signal intensity.

Pointing of the two antennas toward the spacecraft to a small fraction of the primary beamwidth is routinely achieved with the operational antennas of the Apollo Complex. However, this does not guarantee that the peak of the interference pattern generated by two- and three-antenna arrays illuminating the spacecraft will be maximized: the transmitted phases have to be aligned at the spacecraft in order to achieve maximized combined power. Initially, the spacecraft may be located near the peak or the null of the two-antenna interference pattern, but most likely at some intermediate point on the array gain profile. Note that peaks and nulls occur closely spaced in the interference pattern, with fringe distance inversely proportional to the length of the baseline for the two-antenna configurations, requiring precise electronic beam-steering to point the peak of the array pattern towards EPOXI.

During the experiment, the EPOXI high-gain antenna (HGA) was pointed towards Earth: this meant that power levels of only 2 kW had to be used with each of the three antennas, in order not to saturate the CLA readings (even though the transmitters could deliver a maximum of 20 kW). At the beginning of the track, the DSS-25 transmitter was commanded to transmit a 2-kW X-band carrier, and executed an operational uplink frequency sweep to ensure that the SDST aboard EPOXI acquired the signal one-way light-time later

(OWLT = 172 s). One round-trip light-time later (RTLT = 344 s), the DSS-25 ground station receiver acquired and locked onto the downlink carrier and began recording engineering telemetry from the spacecraft. The engineering telemetry contained readouts from the WBAGC, NBAGC, and CLA, providing near-real-time updates of the received uplink carrier power every five seconds.

III. Uplink Array Frequency Predicts

The requirement to maintain the coherence of the uplink array carriers at the spacecraft made it necessary to refine the existing single-antenna predicts, and even develop a new approach to generate greatly improved frequency predicts. With the relatively long baselines formed by the Apollo Complex (258 to 500 m), it was found necessary to refine the positions of the antenna phase centers after position errors as large as 60 cm were discovered in the DSN database. The long baselines lead directly to time-varying differential Doppler between the array antennas (due primarily to Earth rotation), hence, any inaccuracy in the antenna position vectors could lead to significant frequency-prediction errors. In a previous experiment, the position errors in the DSN database led to differential frequency errors of approximately 1 mHz in the ITT frequency predicts ordinarily used in routine tracking operations and also during the previous EPOXI Uplink Array Experiment (February 8, 2008). Accurate position vectors derived from previous very long baseline interferometry (VLBI) measurements were used to yield much more accurate position vectors, resulting in greatly improved frequency predicts for the uplink array.

The pointing-based approach to frequency predicts is based on the observation that the accuracies of the differential Doppler frequencies are the most important components of the predicts for array applications, instead of absolute frequency accuracy. In general, predicts are designed to freeze the received frequency at the spacecraft at a predetermined value, by cancelling Doppler due to Earth rotation and spacecraft trajectory dynamics. However, the SDST aboard EPOXI operates with a 100-Hz loop bandwidth, thus it can easily track out small deviations from the design frequency after signal acquisition. Therefore, it is sufficient to relate the frequency predicts for the array antennas to the reference antenna predicts, which are tracked by the spacecraft.

Referring to Figure 4, the differential position vectors from the reference antenna, DSS-25, to the array antennas (DSS-24/26) have been determined to better than millimeter accuracy using VLBI-derived solutions.[1] These position vectors refer to the phase centers of the antennas, defined as the intersection of the azimuth and elevation axes. The antennas must be pointed towards the spacecraft to within a small fraction of their 70-mdeg beamwidths; however, this is accomplished routinely using operational pointing predicts derived from Earth and spacecraft ephemerides (rms pointing errors for the 34-m Apollo antennas are reported to be approximately 1.5 mdeg [3]).

[1] Chris Jacobs, personal communications, Tracking Systems and Applications Section 335, Jet Propulsion Laboratory, Pasadena, California.

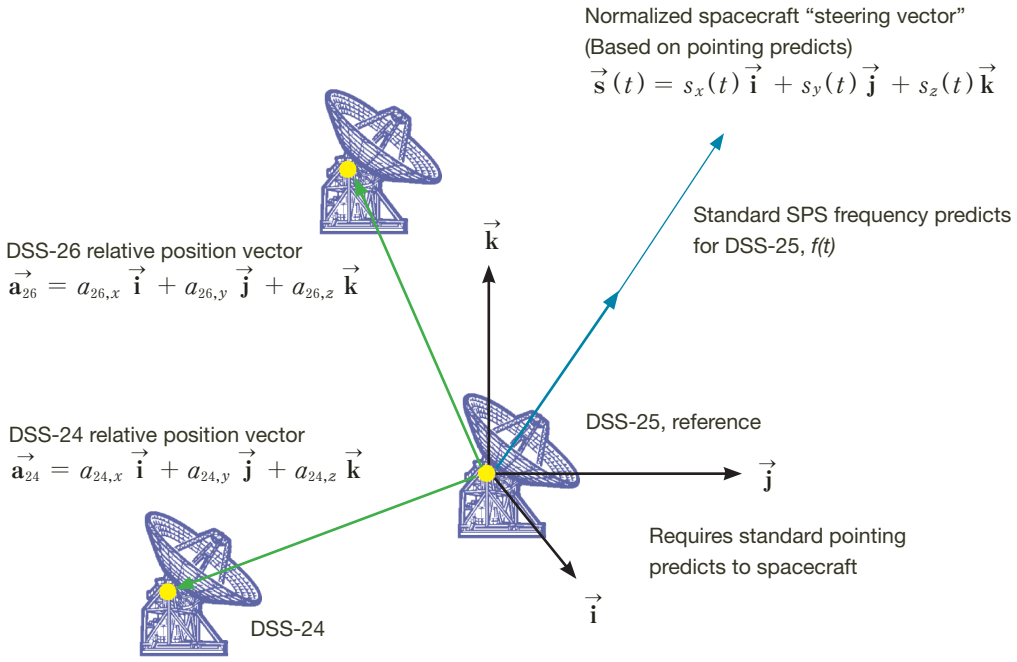

Figure 4. Components of pointing-based frequency predict derivation.

There are two types of pointing predicts employed by the DSN: downlink predicts and uplink predicts. In general, downlink predicts point somewhat behind the current position of the spacecraft in the sky, to account for the change in trajectory during the OWLT between transmission and reception, whereas uplink predicts point ahead of the current spacecraft location to where the spacecraft will be when the transmitted signal from Earth arrives. For this experiment, it was determined that the difference between downlink and uplink predicts was only approximately 1.5 mdeg, which is insignificant compared to the 70-mdeg beamwidth of the 34-m antennas. Therefore, it was decided to use the familiar downlink predicts during this experiment.

Using any adequate single-antenna frequency predict for DSS-25 as reference, the differential frequencies for the auxiliary array antennas can be obtained by forming the inner product of the normalized pointing vector and the position vector from the reference antenna, as shown in Figure 4, for each predict point in time. The single-antenna frequency estimates already contain large Doppler due to relative motion between the reference antenna and the spacecraft. Converting relative velocities to Doppler frequencies at X-band generally requires the use of the relativistic Doppler equation $f = \sqrt{(1 - v/c)/(1 + v/c)}f_C$, where f_C is the carrier frequency in the frame of the reference antenna. The instantaneous pathlength difference to the target between the reference and the auxiliary antennas can be converted to an instantaneous phase difference via multiplication by $2\pi/\lambda$, where $\lambda \cong 4$ cm is the wavelength of the nominally 7.18-GHz carrier. Referring to Figure 4, the instantaneous pathlength difference can be determined as $f_{\Delta,n}(t) = \frac{2\pi}{\lambda}\vec{\mathbf{a}}_n \cdot \frac{\partial}{\partial t}\vec{\mathbf{s}}(t)$ for the n-th auxiliary antenna, n = 24, 26 in the fixed reference frame of the array. The frequency difference between the n-th antenna and the reference antenna in this reference frame

is the rate of change of the pointing vector $p_n(t) = \frac{2\pi}{\lambda}\vec{\mathbf{a}}_n \cdot \vec{\mathbf{s}}(t)$. Finally, the frequency predict for the n-th antenna is the sum of the reference antenna predict and this frequency difference: $f_n(t) = f(t) + f_{\Delta,n}(t)$.

Discrete frequencies can be generated for at most 500 points, which is the limiting size of the operational predict files: hence, each three-hour segment was divided into 490 predict points separated by 22 s. The exciters generate interpolated X-band frequency ramps between predict points, which are then routed to the power amplifiers at each antenna via 16-km optical fibers. Since the rate of change of pathlength difference corresponds to differential velocity, consistent differential velocity estimates can be obtained by differencing the interpolated pathlengths: in this case, pathlengths were computed every 11 s, and used to bracket the desired values at the 22-s spacing. In this way, consistent range and velocity estimates were obtained, meaning that one variable can be derived from the other without significant error.

After the DSN database had been corrected, it was observed that the above procedure was equivalent to and could be carried out more efficiently using Navigation and Ancillary Information Facility (NAIF) range predicts directly [4], together with the improved position vectors, to obtain identical results. Therefore, the actual frequency predicts for DOY-179 used during this experiment were computed using NAIF range predicts every 11 s, and velocities were computed using the differencing method described above. The frequency predicts for the reference antenna, DSS-25, were also derived in this manner for the DOY-179 experiment.

NAIF offers an information system called “SPICE” to assist scientists in planning and interpreting observations from spaceborne instruments. SPICE is also widely used in engineering tasks needing access to space geometry. SPICE is focused on solar system geometry, time, and other related information. The SPICE system includes a large suite of software, mostly in the form of subroutines, which customers use to read SPICE files and to compute derived observation geometry, such as altitude, latitude/longitude, and illumination angles. SPICE data and software may be used within many popular computing environments. The software is offered in Fortran, C, IDL®, and Matlab.

It is important to point out that the NAIF range and velocity predicts were not consistent with each other until recently. Relative range and velocity estimates are obtained through Chebyshev expansion polynomial evaluations, but the relative velocity was not constrained to equal the derivative of the relative position. This inconsistency is not significant for single-antenna applications but manifests itself as a diurnal accumulated phase error (between instantaneous range and velocity predicts) of up to 100 deg for array applications. This has since been corrected in version N62 of the SPICE toolkit.

For comparison during the experiment, and to provide some degree of verification for the pointing-based predicts, three other sets of frequency predicts were derived using different techniques. These were: SPS predicts derived with development software; official ITT predicts using operationally approved software; and custom predicts using a unique technique

whereby "forward" and "backward" predicts were first derived, then averaged to obtain the final result. The key features of these techniques are as follows:

(1) The ITT predicts are generated from modified versions of the NAIF kernels, which take into account general relativity effects such as pathlength increases from gravitational effects of massive bodies. Tropospheric compensation based on a seasonal variation model is supported by both ITT and SPS but was not used in this experiment.

(2) For the custom predicts, SPICE routines from NAIF were used in developing a Fortran program to provide the various parameters of interest in support of the EPOXI experiment. Directions (AZ-EL), geometric pathlengths, RF phases, and Doppler frequency shifts for DSN stations to EPOXI were evaluated. In implementing the program, a forward pass from the transmit station to the spacecraft, and a backward pass from the spacecraft back to the station is first computed, in order to calculate the transmit frequency F_t at the given time T_0 at the transmit station needed to provide the given frequency F_r at the spacecraft at calculated time T_r. The actual values are obtained by averaging over the forward and backward passes. In this way, a set of ephemeris data at the desired range of UTC times is calculated. The main NAIF routine utilized for these calculations was "SPKEZR," which returns the state (position and velocity) of a target body relative to an observing body, optionally corrected for light time (planetary aberration) and stellar aberrations.

(3) SPS predicts used the same algorithms as the ITT predicts, except that operational constraints were relaxed, so that the most recent updates of ephemerides and kernels are used; which, however, may not be approved for operational applications.

Comparison of the three techniques (ITT, SPS, and custom) to the pointing-based predicts are shown in Figures 5(a–c) for the third predict interval spanning 1500–1800 UTC on DOY-179. The blue curves refer to the DSS-24/25 baseline, and the red curves refer to DSS-25/26. Note the excellent agreement between the pointing-based and the custom predicts: the peak difference is about 0.5 deg for DSS-24/25 and about 2.25 deg for DSS-25/26. These differences would not lead to any measurable power fluctuations at EPOXI, hence, these predicts would yield essentially identical performance during the third interval.

However, both the SPS and ITT predicts accumulate significant phase linearly with respect to the pointing-based predicts, amounting to approximately –17 deg and –25 deg for ITT and SPS, respectively, along the DSS-24/25 baseline, and 20 deg and 32 deg for the DSS-25/26 baseline. It can be inferred that the developmental SPS and operational ITT predicts differ from each other by only –8 deg and 12 deg for the two baselines, hence, either predict set would yield very similar performance during the experiment. These differences can be attributed to different kernels used, and to different constraints imposed on the solution. It is believed that the pointing-based predicts provide the most accurate predicts for arraying applications, as these were developed specifically to minimize differential phase errors and provide consistent range and velocity predicts, which eliminates the accumulated differential phase error due to diurnal variations.

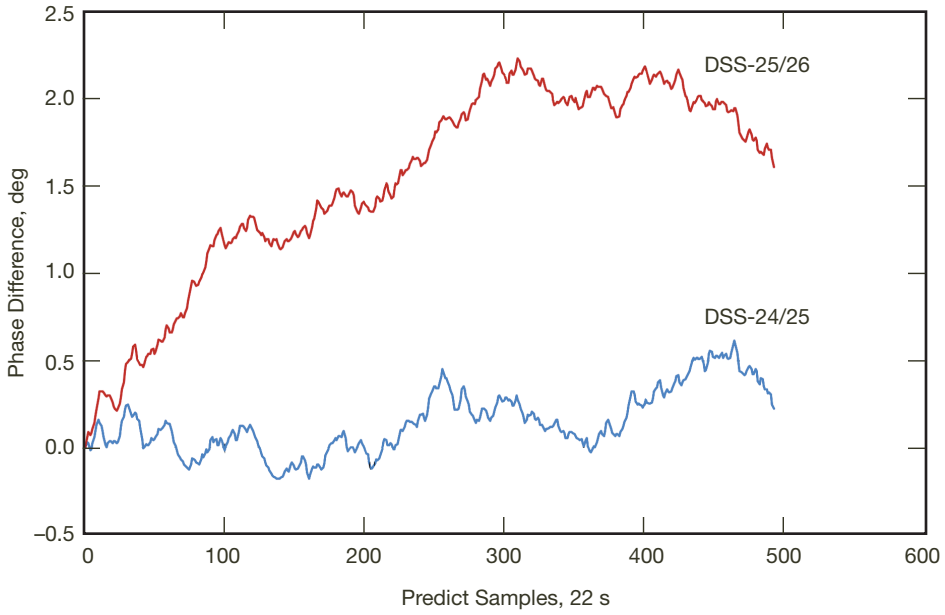

Figure 5(a). Comparison of differential phase errors for pointing-based and custom frequency predicts: third tracking interval, 1500–1800 UTC. Maximum observed difference was 0.5 deg for DSS-24/25, 2.2 deg for DSS-25/26.

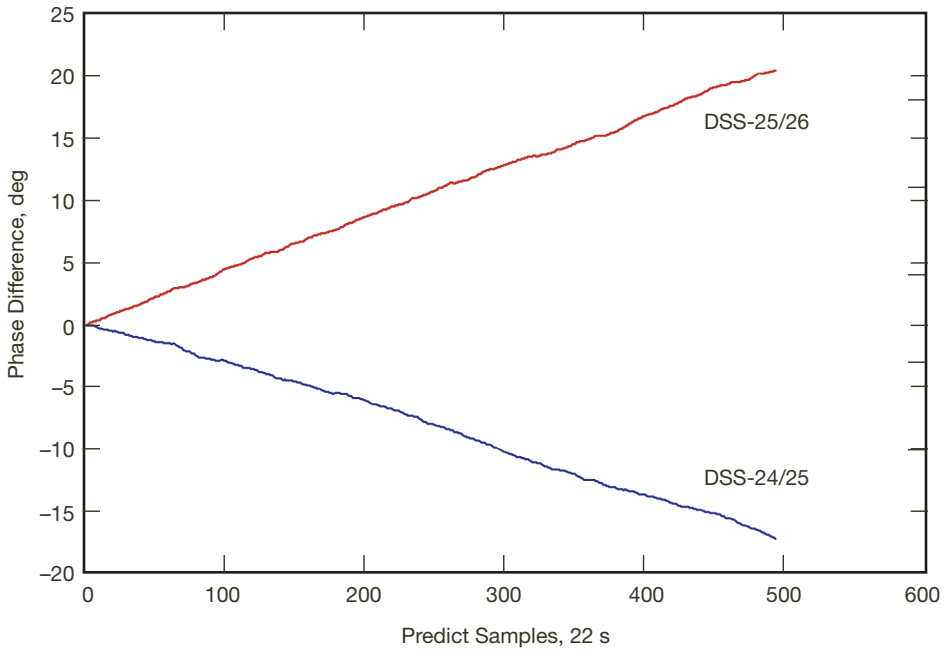

Figure 5(b). Comparison of differential phase errors for pointing-based vs. ITT frequency predicts: third tracking interval, 1500–1800 UTC. Maximum observed difference was –17 deg for DSS-24/25, 21 deg for DSS-25/26.

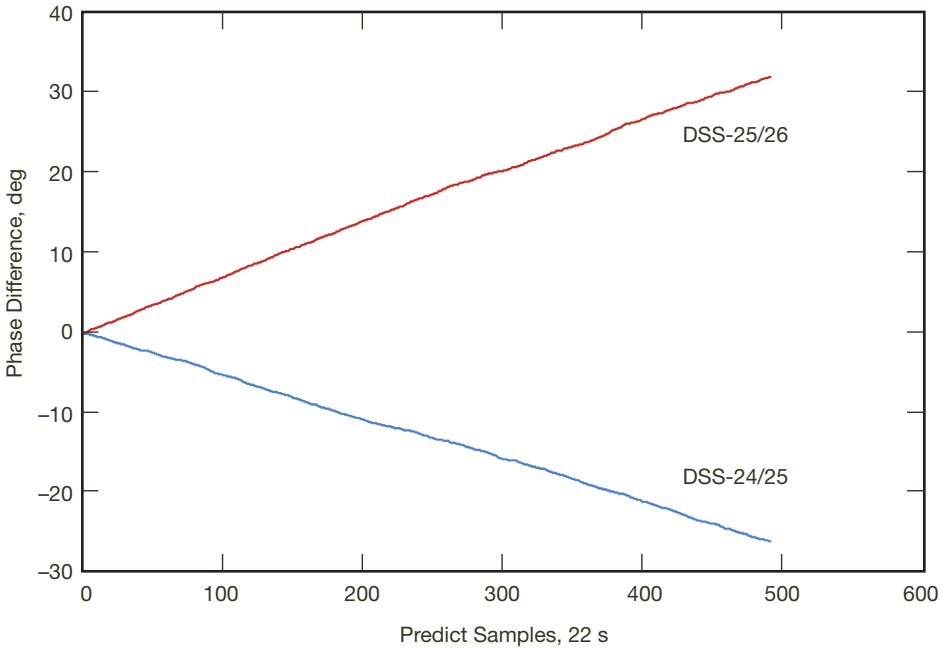

Figure 5(c). Comparison of differential phase errors for pointing-based vs. SPS frequency predicts: third tracking interval, 1500–1800 UTC. Maximum observed difference was –26 deg for DSS-24/25, 32 deg for DSS-25/26.

IV. Ground System Monitoring and Control

Slowly varying phase drifts in the ground system are monitored and controlled by the Phase Comparator and Control Assembly (PCCA), located in a half-rack at SPC-10. The PCCA contains a signal distribution assembly (SDA), two Phase Comparator Assemblies (PCAs) to measure round-trip and cross-phase from all three Apollo antennas, and a Phase Modulation Assembly (PMA) that can be used to add correction phases to either the DSS-24 or the DSS-26 carriers. A block diagram of the PCCA is shown in Figure 6.

The theory of operation for the PCCA has been fully described in a previous article [5]. In summary, the ground system consists of the X-band exciters at SPC-10, X-band couplers and a "round-trip" Phase Comparator Assembly at SPC-10, optical fibers for signal distribution to the transmitting antennas, X-band couplers at the output of the power amplifiers at each antenna, and additional optical fibers to return the coupled signal samples to SPC-10 for comparison. The two-way optical fiber distribution network to and from the antennas is located in the same bundle for most of the 16-km distance from SPC-10 to the Apollo cluster, ensuring similar thermal behavior for the outgoing and returning signals. At the Apollo Station, the individual fibers are broken out from the common bundle and routed to their respective antennas, typically a distance of a few hundred meters, over which the fibers may experience independent thermal environments.

Following power amplification at the antenna pedestal room, a small fraction (–54 dB) of the amplified X-band signal is coupled off the waveguide and routed back to SPC-10, where the phase of this "sample" signal is compared to the transmitted phase using the real-time Phase Comparator Assembly (PCA). The inputs to the PCA are the outgoing (reference) and

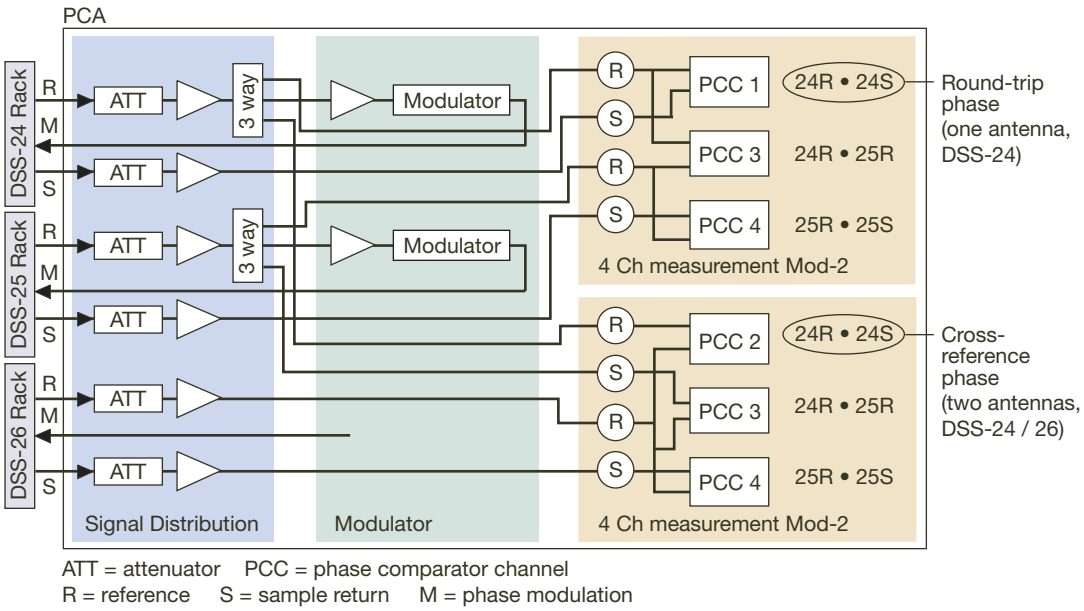

Figure 6. Block diagram of Phase Comparator and Control Assembly, including the SDA, PMA, and PCA.

return (sample) signals from the Apollo antennas. The PCA outputs consist of complex samples of equal magnitude representing the phase difference between the reference and sample of each antenna, called the "round-trip" phase, which is used to measure changes in total pathlength, or the difference phase between the reference outputs of two different antennas at SPC-10. For example, the upper right output in Figure 6 represents the round-trip phase of DSS-24, whereas the lower right output represents the cross-reference phase between DSS-24 and DSS-26.

The PCA was designed to provide at least 5 deg of phase accuracy after calibration. The power levels of the reference and sample signals were measured and adjusted at Goldstone using a transmitter setting of 2 kW, which is the lowest recommended power level with these 20-kW transmitters, without risking possible transmitter instabilities. A 100-Hz low-pass filter is applied at baseband after the RF mixers to reduce noise and smooth out fluctuations. After amplification, the baseband signals are sampled by an 8-channel, 16-bit analog-to-digital converter. These samples are processed in Labview to remove any DC offsets and gain differences between the I and Q channels (causing ellipticity in the I-Q plots), and the appropriate phases computed as the 4-quadrant arctangent of the I and Q baseband signals.

The host computer of the PCCA was used to measure the round-trip and cross-phases between the two baselines, namely the DSS-24/25 and the DSS-25/26 baselines. Baseband I and Q samples were taken 100 times per second, and averaged to obtain 1-s estimates from which the phases were computed. A complete record of the cross-phase behavior of the DSS-24/25 baseline for the entire track is shown in Figure 7(a), for all three predict sets. Note that the cross-phase jumps to different random values each time a new predict set is loaded, since the exciters were not designed to maintain constant phase between predict sets. However, the cross-phase remained essentially constant throughout each predict set, indicating excellent long-term phase stability. The slow phase drift from ~130 to ~170 deg during the

third predict interval is believed to be an artifact due to accumulated timing error in the host computer, since a corresponding drop in combined power was not observed in the EPOXI AGC readings. The computer timing has since been improved, and will be tested during subsequent experiments.

Figure 7(b) represents the difference of the round-trip phase measured for the DSS-24/25 baseline. Note that the round-trip phase is not subject to phase jumps when a new predict set is loaded in, since both the return signal and transmitted carrier pick up the same phase

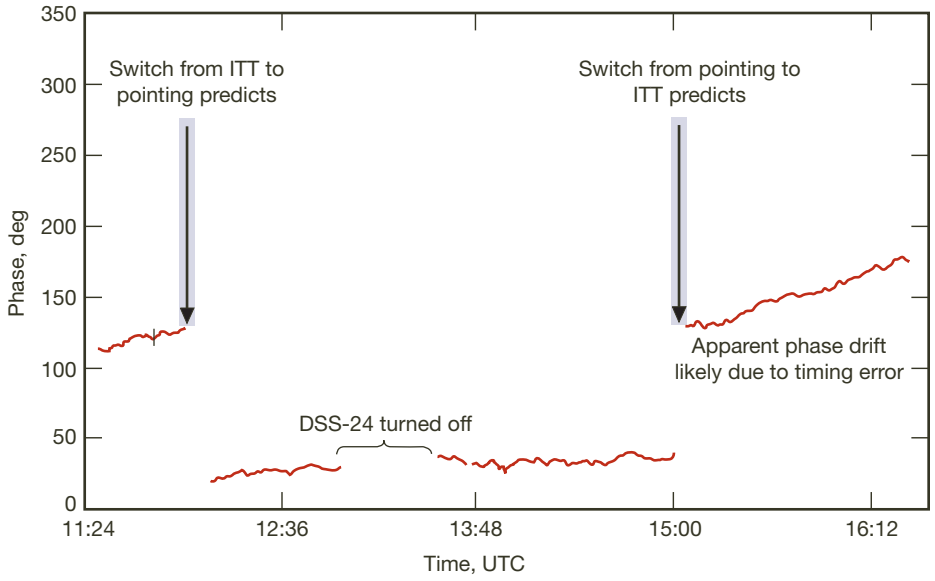

Figure 7(a). DSS-24/25 cross-phase history covering the entire track.

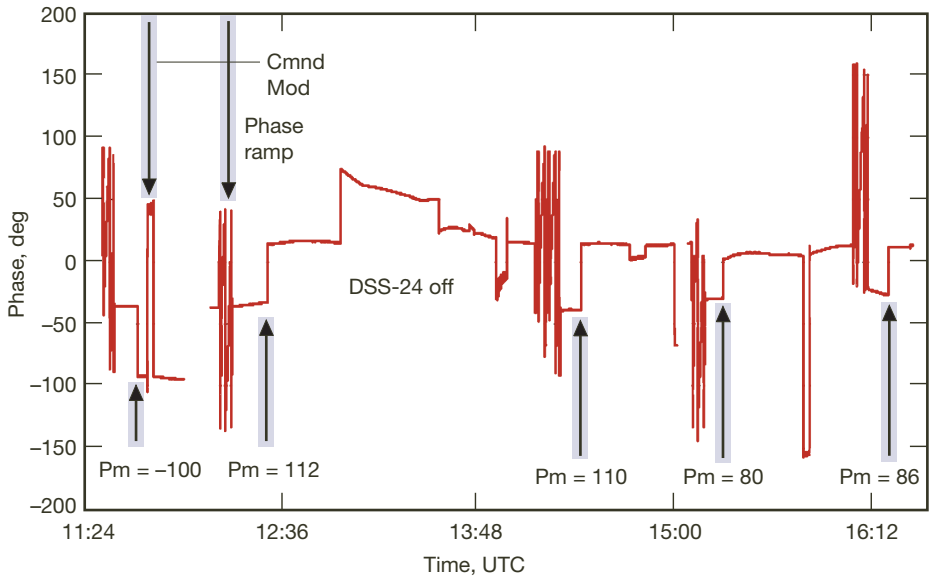

Figure 7(b). Round-trip phase history: DSS-24 minus DSS-25 phase divided by two (accounting for doubling due to round-trip from the antenna back to SPC-10). Examples of command modulation (Cmnd Mod), phase modulation (Pm), and phase ramp signatures are indicated.

shift: only changes in phase applied after the reference coupler are registered by the PCA, whether they are due to an applied phase offset via the PMA, or to unintended change in electrical pathlength due to thermal effects. It can be seen that the round-trip phase remained stable throughout the track, except when command modulation or phase ramps were applied, and except for approximately half an hour around 14:00 UTC when the DSS-24 power was shut off, hence, the phase trajectory during this time represents DSS-25 pathlength change, which can be much greater than the differential phase between the two antennas. Therefore, no phase corrections had to be applied to compensate for ground-phase drift during this experiment. The DSS-25/26 baseline exhibited similar phase behavior throughout the EPOXI track.

V. EPOXI Uplink Array Experiment Results

The EPOXI track started at 11:00:00 UTC on DOY-179 (June 27, 2008), when the spacecraft was at an elevation of approximately 30 deg, rising. Operational ITT frequency predicts were used for the first predict interval. These predicts were not corrected for tropospheric variations; however, differential phase build-up due to the troposphere was predicted to be insignificant above 30-deg elevation, leading at most to 20 deg over the DSS-24/25 baseline. The DSS-25 transmitter was configured for 2 kW of uplink power, with right-hand circular polarization (RCP). Transmission started shortly after 11:00 UTC, with a predetermined operational frequency ramp sequence to ensure carrier capture by the SDST phase-lock loop. The OWLT during this experiment was approximately 172 s, or 2 min 52 s, hence, the first downlink power measurement was observed nearly 6 min after transmission. Data were recorded with the WBAGC, NBAGC, and the CLA. As shown in Figures 9, 10, and 11(a–d), the WBAGC and CLA records were received shortly after 11:00 on the spacecraft, one light-time after transmission started. In these and subsequent figures, the time axis refers to spacecraft time in UTC, not ground time. The SDST locked onto the DSS-25 signal at roughly –103 dBm of carrier power, close to the predicted value. This power level was well within the linear region of the WBAGC, whose resolution is 1.5 dB as shown in Figure 9, but near the saturation region of the CLA; therefore, the readings in Figure 10 appear lower than the true power before correction.

Corrections were applied to the digital numbers (DN) returned by the CLA during post-processing, via the temperature-dependent distortion profiles shown in Figure 8(a). The calibration curves refer to 25 deg C and 60 deg C baseplate temperatures; however, the true baseplate temperature during this experiment was approximately 42.5 deg C, hence, the corrections were interpolated to this temperature. The corrected values were used to generate Figures 11(a–d), all of which are referenced to spacecraft time (termed “spacecraft event time” or SCET.) As an example of the use of these correction curves, a DN reading of 3500 yields –106 dBm, as shown in Figure 8(a).

After the SDST locked up to the DSS-25 carrier, 16-kHz uplink modulation was turned on, and a sequence of 10 No-Op commands was sent to the spacecraft at about 11:12 SCET, recorded as a drop in CLA power from –102.7 dBm to –113 dBm in Figure 11(a). Since the CLA measures the carrier power, the power in the sidebands is lost to the CLA. However, since the modulation index was 0.94 rad, only 5.8 dB should have been lost to modulation,

not the approximately 10-dB drop shown in Figure 11(a). One reason for the discrepancy is that the correction curve is different for the modulated carrier, but this correction curve was only available for 25 deg C, hence it could not be interpolated to a different baseplate temperature. As an example of correction for the modulated carrier, it can be seen from Figure 8(b) that a DN reading of 4400 yields a modulation-corrected value of –95.5 dBm at a temperature of 25 deg C.

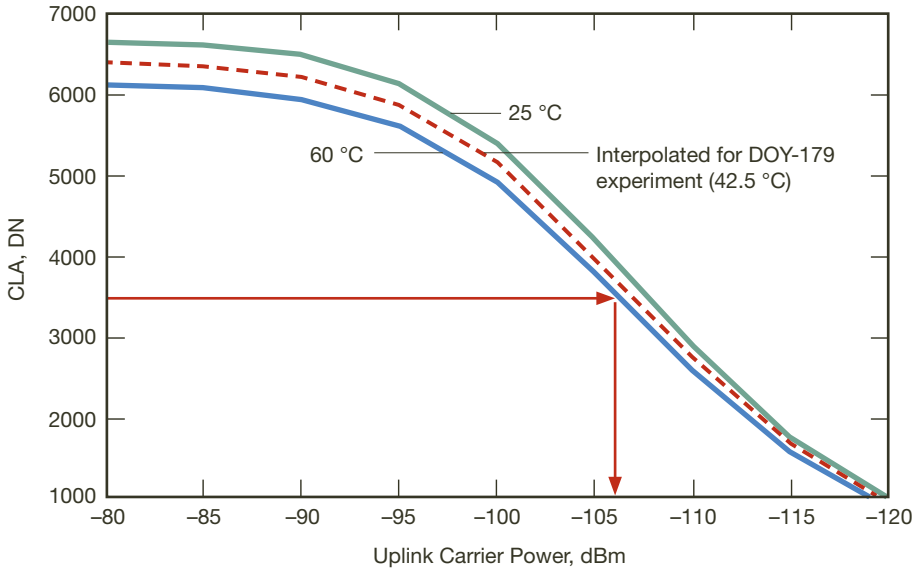

Figure 8(a). CLA correction curve from DN to dBm, carrier only, for the two calibrated baseplate temperatures of 25 deg C and 60 deg C, as well as the interpolated curve for 42.5 deg C.

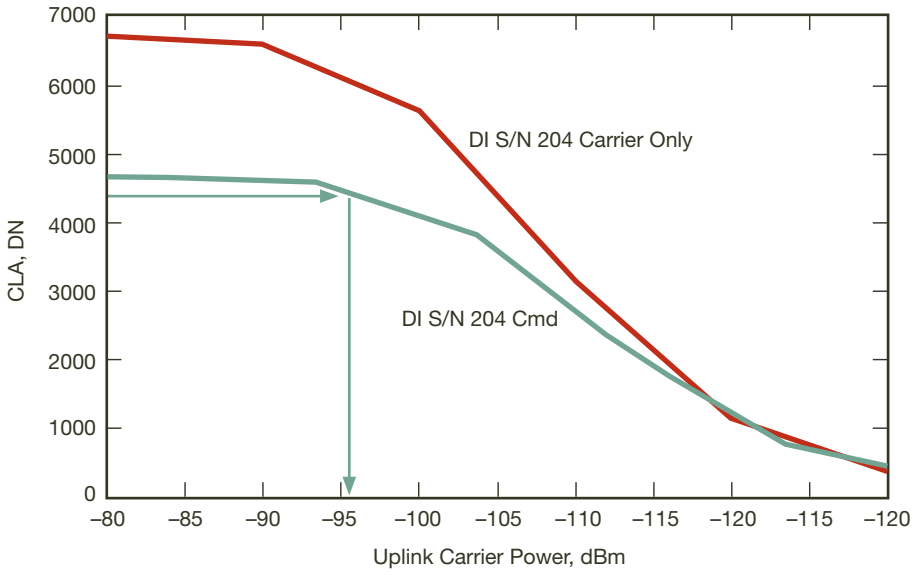

Figure 8(b). CLA correction curve from DN to dBm, carrier only plus carrier with command on, at a calibrated baseplate temperature of 25 deg C.

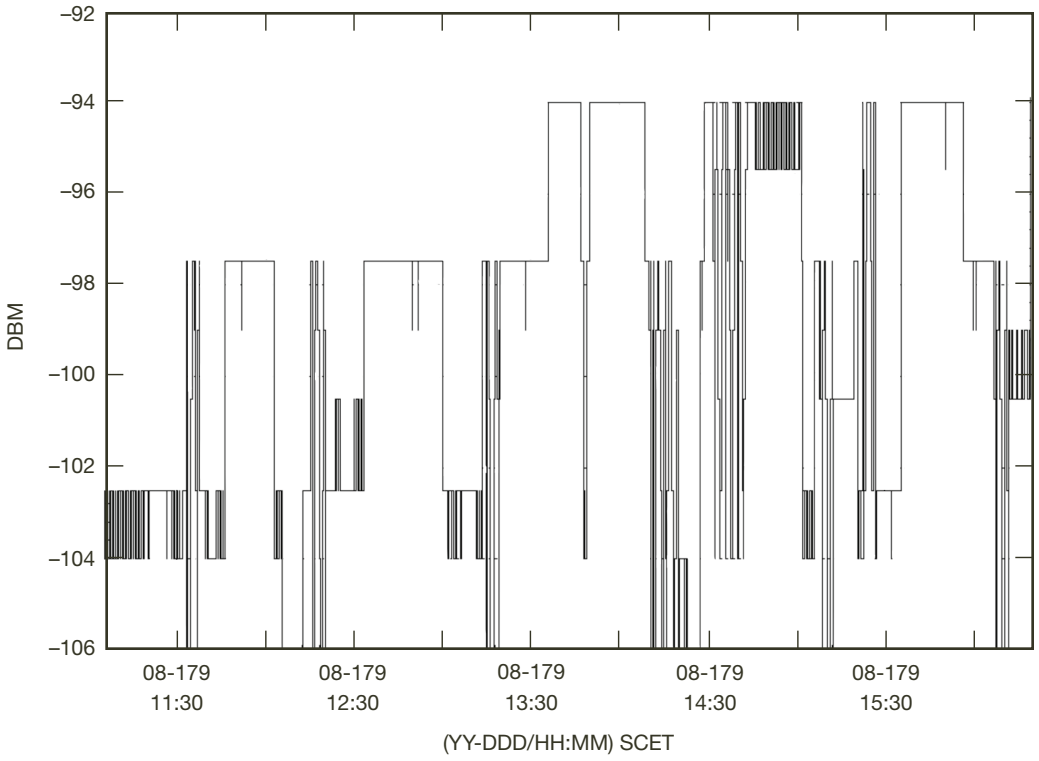

Figure 9. Recorded WBAGC data, obtained post-track from the telemetry delivery subsystem.

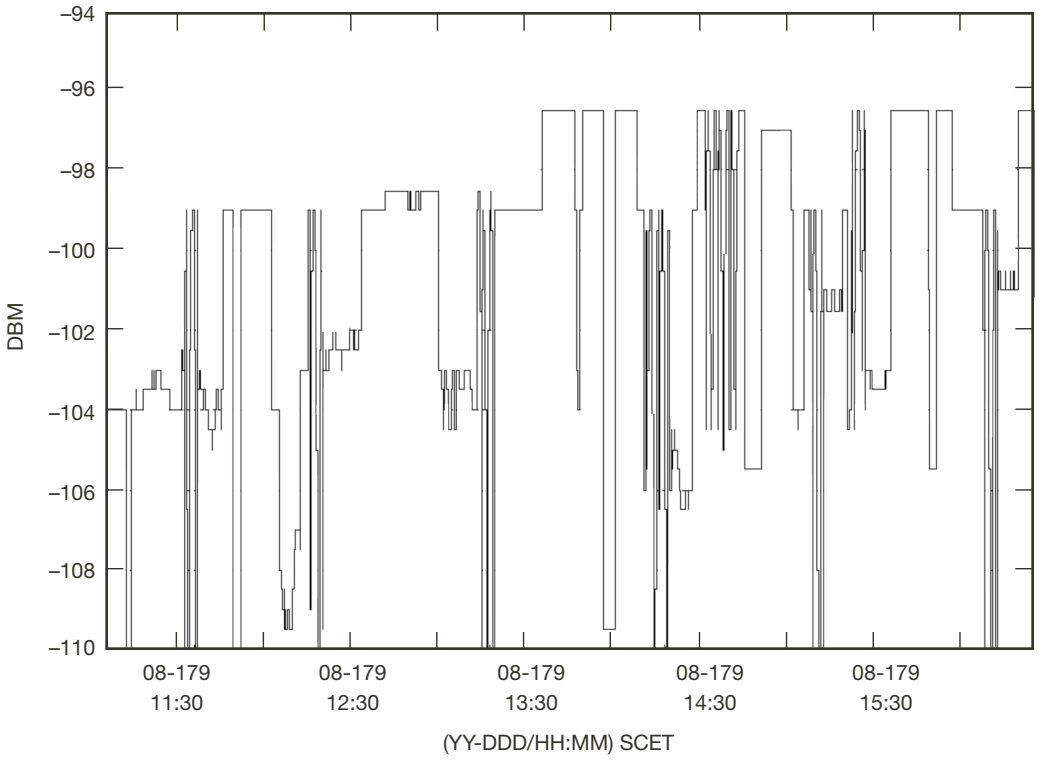

Figure 10. Recorded CLA SNR data, obtained post-track from the telemetry delivery subsystem.

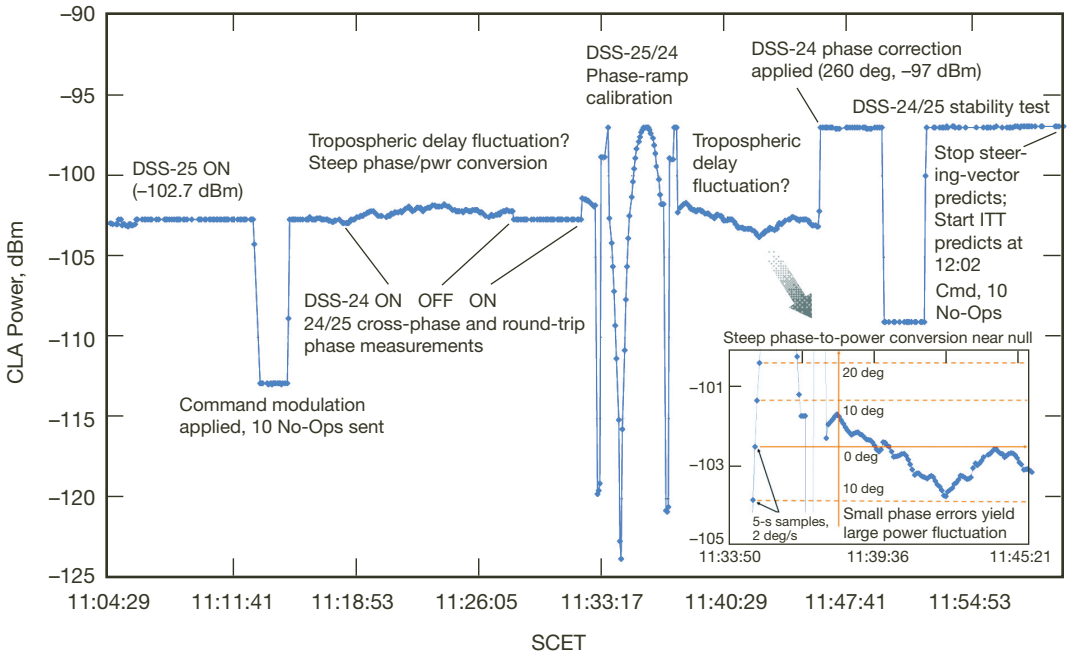

Figure 11(a). Real-time CLA data at the start of the experiment, including examples of No-Op commands and phase-ramp sequence for estimating optimum phase.

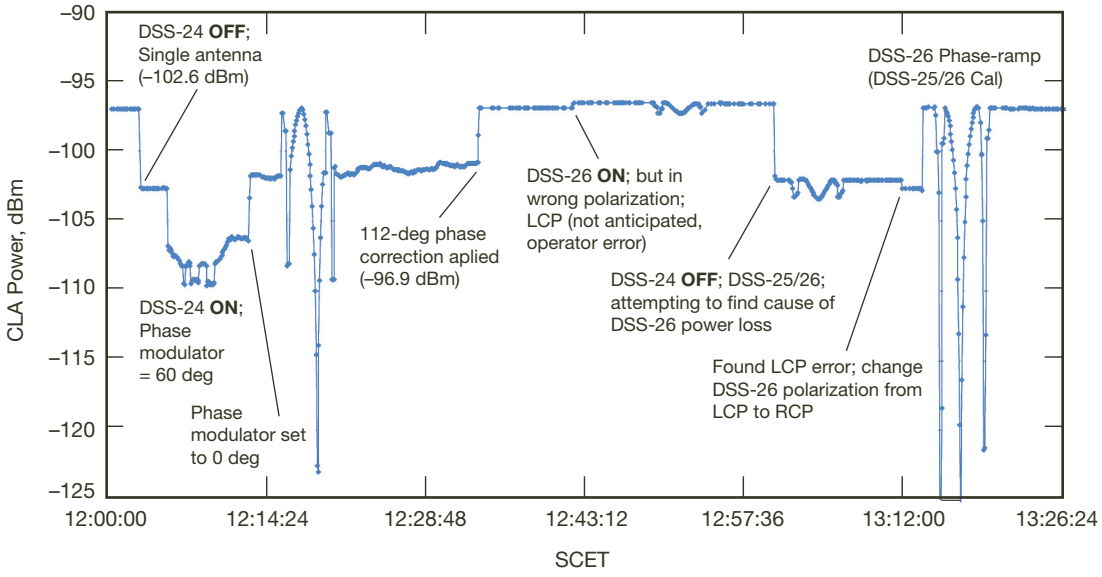

Figure 11(b). Real-time CLA data, first part of second predict set.

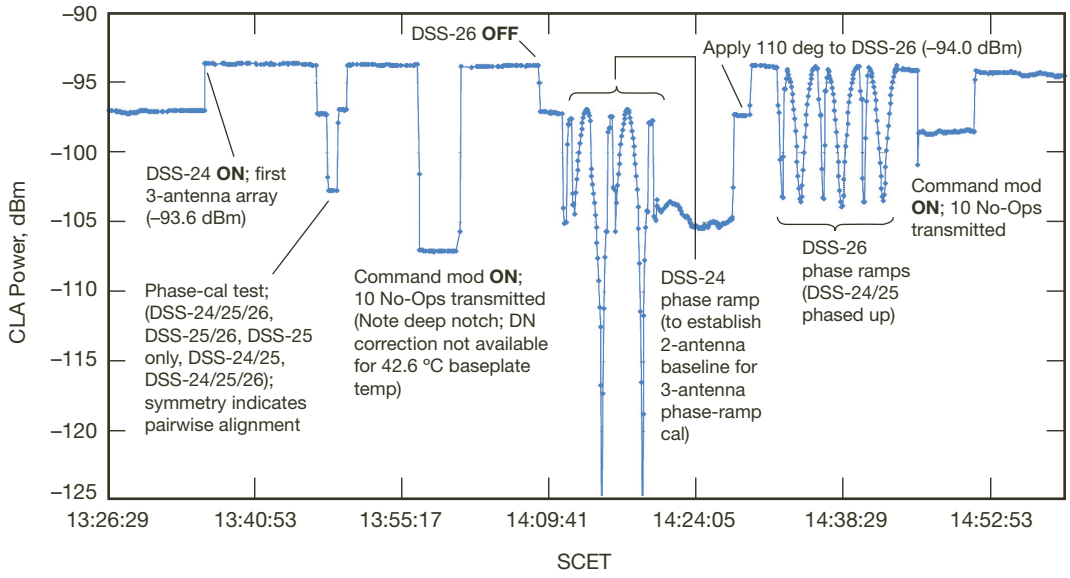

Figure 11(c). Real-time CLA data, second part of second predict set.

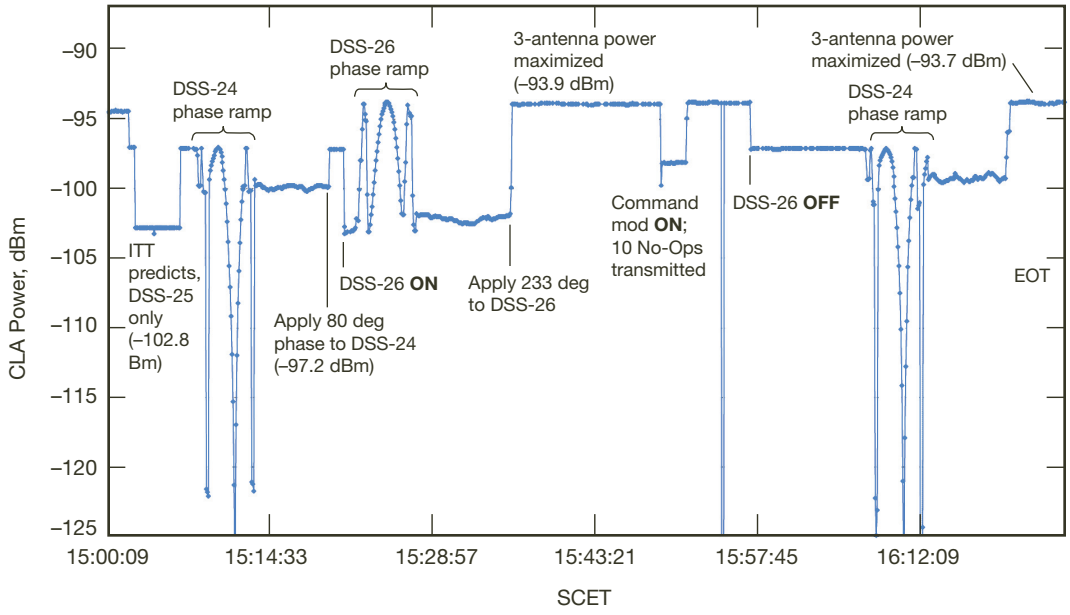

Figure 11(d). Real-time CLA data, last part of experiment, third predict set.

Following successful reception and acknowledgement of all 10 of the first set of No-Op commands by EPOXI at about 11:13 SCET, the modulation on the DSS-25 carrier was turned off, the DSS-24 transmitter drive was turned on, and the sum of the two carriers was received by EPOXI at 11:18 SCET, as shown in Figure 11(a). The AGC readings were updated every 5 s, with a DN resolution of 0.3 dB on the CLA, 0.5 dB on the NBAGC, and 1.5 dB on the WBAGC.

The far-field intensity produced by the two-antenna uplink array can be steered over the spacecraft electronically from the ground by varying one of the transmitter phases. When the differential phase varies over the entire range of $(0, 2\pi)$ rad, it is guaranteed that the peak of the intensity pattern illuminates the spacecraft, provided the phases are stable over the duration of the sweep. This power variation can be measured by the spacecraft AGC to monitor the instantaneous carrier power at the spacecraft. The results of these measurements are then relayed to the ground as engineering data, and evaluated to determine the optimal phase adjustment required to phase up the signals at the spacecraft.

Phase ramping was initiated at 11:33 SCET via the digitally controlled phase modulator on the DSS-24 channel, in order to determine the optimum phase. This technique, termed "Nearby Spacecraft" calibration, was conducted on the DSS-24/25 baseline by ramping the DSS-24 phase from 0 to 2π rad at a rate of 2 deg per second using the PMA, and the downlinked spacecraft AGC power readings were monitored in real-time at SPC-10. As can be seen in Figure 11(a), this resulted in an increase in measured array power from –102.7 dBm to –97 dBm, corresponding to a two-antenna array gain of 5.7 dB, close to the theoretically predicted 6-dB gain. Each phase ramp consists of a preamble, which is a pattern of 0, 90, 180, 270 deg applied for 10 s each to mark the beginning of the ramp sequence, followed by a linear phase ramp of 2 deg per second applied for 180 s, concluding with a repeat of the preamble pattern to help determine the end of the measurement sequence. The phase ramp itself always starts and ends at 0 deg, to simplify the calibration process. The phase corresponding to maximum combined power is the desired calibration phase, which was determined to be 260 degrees for the first phase ramp on the DSS-25/24 baseline. After applying the optimum phase, the combined power immediately jumped to –97 dBm at 11:45 SCET, verifying two-antenna power maximization.

Note that around 11:40 SCET, the combined power tends to fluctuate when the phase difference between the two antenna carriers is large, placing the combined signal on a steep part of the phase/power curve. On this part of the curve, the two-antenna array power changes rapidly with small phase variations, hence, the array serves as a very sensitive instrument for measuring tropospheric fluctuations. As described in the inset in Figure 11(a), on the part of the power/phase curve corresponding to approximately 100 deg of differential phase error, the combined power fluctuates near the single-antenna power level –103 dBm, likely due to differential tropospheric delay fluctuations on the order of 1 mm over the 258-m baseline.

After verification of two-antenna combined power magnitude and stability, the first sequence of 10 "arrayed" No-Op commands was transmitted to EPOXI by the two-antenna array at approximately 11:50 SCET. Since the uplink subcarrier is 16 kHz, corresponding to

a wavelength of 18.75 km, the maximum differential delay between any two antennas in the Apollo Complex never exceeds 0.5 km/18.75 km = 0.027 wavelengths, thus it can be neglected: this means that delay compensation is not required (0.5 km is the maximum extent of the array, along the DSS-24/26 baseline). Therefore, the modulation was distributed to all three antennas via an active three-way splitter/amplifier signal distribution assembly constructed especially for this purpose. The calibration curve for the EPOXI SDST with command "ON" is significantly different than with command "OFF" at high power levels (towards the left side of the curve), as can be seen in Figure 8(b). One of the reasons for the greater than expected 5.8-dB drop in carrier power recorded by the CLA with modulation "ON" is that the calibration curve was only available at a single temperature (25 deg C), hence, it could not be interpolated to the much higher operating baseplate temperature of 42.5 deg C to obtain better correction. Another contribution could be from the test modulation input ports, some of which have not been calibrated, and could have altered the modulation indices significantly from their design value. Finally, it was noticed towards the end of the experiment that by changing the distribution cables, some of these problems could be remedied, hence, the last two No-Op commands of the track show the expected 6-dB drop in carrier power as measured by the CLA, indicated in Figures 11(c) and 11(d).

The first ITT predict set ended at 12:00:00 UTC, and approximately 3 min later at the spacecraft due to the OWLT. For the second tracking interval, the pointing-based frequency predicts were selected; these predicts were derived specifically for array applications and hence yielded consistent range and velocity estimates. The events of the second predict interval are depicted in Figures 11(b) and (c). After the predicts were started, a phase-ramp calibration was initiated at 12:15 SCET, from which the optimum offset was determined to be 112 deg, which indeed maximized the combined power at EPOXI at –96.9 dBm after application via the PMA. Next, the DSS-26 transmitter was turned on, and the signal reached EPOXI just before 12:43 SCET. The second phase ramp on DSS-26 was initiated in order to calibrate the three-antenna array; however, a total power variation of only about 1 dB was observed, instead of the expected 9.5 dB. Further tests led to the realization that the polarization of DSS-26 was accidentally set to LCP instead of RCP, causing large losses in received power from this station.

To help determine the cause of the power drop, the DSS-24 transmitter was turned off at approximately 13:00 SCET, shown in Figure 11(b), and a phase ramp applied to the DSS-26 PMA, which also yielded small power variations. After ruling out transmitter power anomalies and antenna pointing errors, it was postulated that DSS-26 may have been set up in the wrong polarization (LCP instead of RCP): this was indeed found to be the case.

After the DSS-26 polarization was reset to RCP, another two-antenna phase ramp was initiated on the DSS-25/26 baseline, and the expected two-antenna response observed starting at 13:15 SCET. Following calibration along this baseline, the DSS-24 transmitter was turned back on, whereupon the combined three-antenna power jumped to –93.6 dBm at about 13:36 SCET, close to the theoretically expected combined power of –93.3 dBm. Since the DSS-24/25 baseline was already calibrated and the frequency predicts maintained constant phase at the spacecraft, this two-antenna pairwise calibration successfully completed the three-antenna calibration process.

Successful three-antenna calibration via the pairwise method was followed by a phase symmetry test at 13:48 SCET. At first, DSS-24 power was turned off, followed by DSS-26 to measure the single-antenna power of DSS-25, then the power of each was turned on in the opposite order. The slight asymmetry of the shoulders indicates somewhat better phase alignment along the DSS-24/25 baseline than along the DSS-25/26 baseline, which could account for some array gain loss; however, it was not deemed significant. Another set of 10 No-Op commands was transmitted at 13:56 SCET, all of which were successfully received.

Starting at about 14:10 SCET, the DSS-24/25 baseline was recalibrated to reduce any small phase drifts since the last calibration, in order to initiate the three-antenna calibration process. Here, a calibrated pair is viewed as a single antenna, and the third antenna is phase-ramped to maximize array power. Since a calibrated baseline consists of the sum of two carrier amplitudes, but phase ramping of the third antenna can only cancel one of the amplitudes, the combined power varies between the three-antenna and single-antenna powers: this type of phase-ramp pattern can be seen in Figure 11(c), starting at about 14:30 SCET. Another set of 10 No-Op commands was transmitted at 14:45 SCET, before the end of the second frequency-predict set at 15:00 UTC.

For the third and final predict interval, the ITT predicts were loaded in once again to obtain similar total data volumes for comparison (three hours of data for the pointing-based predicts, two and a half hours total for the ITT predicts), and to facilitate more accurate comparison of data both at the beginning and the end of the long second predict interval. In the third predict interval, the three-antenna calibration was repeated, three-antenna stability data were collected, and one more set of No-Op commands was transmitted successfully to EPOXI.

V. Future Directions in Uplink Array Technology Development

The EPOXI Uplink Array Experiment of June 27, 2008, successfully demonstrated the following techniques critical to uplink arraying, in collaboration with a real deep-space probe (EPOXI), and in a realistic operational environment:

(1) Two- and three-element arrays of operational 34–m antennas with 5-mdeg array beamwidth have been calibrated using a nearby spacecraft (EPOXI) to maximize received array power.

(2) Array phase alignment was maintained and the spacecraft was successfully tracked over a large (80-degree) portion of the sky for more than five hours.

(3) Both CW and realistic No-Op commands were transmitted at the maximum rate of 2 kbps, all of which were successfully received and acknowledged by EPOXI, proving operational uplink array command capability.

It remains to be shown that the uplink array can be calibrated independently of the target spacecraft, using the “Moon-bounce” calibration technique described in [6,7]. Preliminary tests of the Moon-bounce approach indicate that calibration should be feasible to approxi-

mately 10 deg of phase for a two-antenna array. Extremely small array losses of 0.3 to 0.4 dB were observed for two- and three-antenna arrays, respectively, consistent with the analytical predictions for a high-quality operational 34-m antenna array.

For the two- and three-element arrays used in this experiment, the combining losses were due to a combination of ground-system differential phase-drift, small residual frequency-predict errors, tropospheric delay fluctuations over the array, and blind-pointing errors. The magnitudes of these error sources, and their contribution to the total phase error, are as follows:

(1) Ground distribution system, which now contributes ±10 deg of differential phase error, but can be improved to ±5 deg with an automated closed-loop upgrade to the PCCA.

(2) Tropospheric delay fluctuations of about ±10 deg above 30 deg elevation.

(3) Frequency predicts, differential phase errors: on the order of ±5 deg over 2 hours.

(4) With 34-m antennas, the beamwidth is about 70 mdeg, and blind-pointing accuracy is on the order of 1.5 mdeg rms under good conditions (from [3], wind speed of 10 mph, 3 mdeg for 95 percentile corresponding to 2-sigma).

Since these error sources are independent, we can root-sum-square the phase errors (assuming ±5 deg for closed-loop ground control), and obtain a total phase error of approximately 12 deg rms, which, together with the 1.5 mdeg rms pointing error, yields about 0.2 to 0.3 dB loss for a two-antenna 34-m array, consistent with the losses observed during the recent EPOXI experiment.

Water Vapor Radiometers for Real-Time Tracking of Tropospheric Phase Fluctuations. An inherent advantage of using only a few 34-m antennas (instead of a much larger number of small antennas) is that we can envision placing a water-vapor radiometer next to (or on) each antenna to estimate the delay variations due to tropospheric fluctuations in real time, and use these estimates to compensate for differential delays. This approach could reduce the contribution of the tropospheric delay variations from ±10 to approximately ±5 deg, resulting in rms phase error of only 8.7 degrees. This approach would further reduce array loss, possibly to as little as 0.1 to 0.2 dB when using high-quality 34-m antennas.

Emergency Communications with Uplink Arrays. The simplest type of spacecraft emergency occurs when the ephemerides of the spacecraft are known, but unintended mispointing of the HGA has disrupted the link, and the spacecraft is too far to rely on the low-gain antenna (LGA) to re-establish communications using only a single uplink antenna. In this case, the increased effective isotropic radiated power (EIRP) due to coherent arraying can be used to advantage in re-establishing the command and communication links.

A more complicated emergency situation occurs when communications have been disrupted long enough to render the spacecraft ephemerides inaccurate, so that predict-driven array pointing is no longer feasible. In this case, some form of spatial scanning is required to locate the spacecraft and re-establish the link. For an array composed of a large number

N of small antennas, the solid angle subtended by the beamwidth of the N-element uplink arrays is a factor of N smaller than the primary elemental beam for a maximally compact array; for widely separated elements, the array beamwidth is correspondingly smaller. However, the peak EIRP of the array is greater by a factor of N^2 than the EIRP of the individual elements, hence, the dwell-time required to achieve a given probability of detection at the spacecraft is reduced by a factor of N^2, but now N array beamwidths have to be scanned. Therefore, the time required to scan an uncertainty region corresponding to the primary beam with a maximally compact array is reduced by a factor of N over the time required for a single antenna, when the antenna powers are fixed. But this is only a limiting-case bound, since shadowing and other practical considerations require separation of the array elements, spoiling the assumption of a maximally compact array. Hence, in a practical operational scenario, scanning of an uncertainty region equal to the primary beam with an uplink array may take more or less time than with a single antenna, depending on array geometry and on the number of array elements.

The conclusions regarding emergency communications are significantly different with an array composed of a few large high-quality antennas, such as the 34-m antennas of the Apollo cluster. In this case, the regular structure of the far-field array pattern (as shown in Figure 3) can be used to reduce the time it takes to locate a lost probe. In effect, a very efficient "array search" can be performed simply by adding 180 deg to the phase of one of the antennas, thus exchanging the location of the array peaks and nulls within the entire primary beam. For two antennas, complete coverage of a primary beamwidth using this phase-switching operation requires two units of dwell-time, where each dwell-time is designed to achieve a given probability of detection at the spacecraft. However, since the EIRP is greater by a factor of 4, the net result is that a two-element array is able to scan the space corresponding to the primary beam a factor of 2 faster than a single antenna.

If the uncertainty region exceeds the primary antenna beamwidth, then the conventional array-search technique employing single-antenna beam-shifting operations must be used to cover the large (multiple primary-beam) uncertainty region: however, the two-antenna array still benefits from the factor of 2 reduction in time needed to search each primary beam.

Uplink Array Applications to Planetary Radar. The application of uplink arrays to planetary radar is attractive primarily because the greater EIRP of the array tends to mitigate the inverse R^4 loss suffered by radar signals, which severely limits the effective reach of radar for planetary applications. However, there are additional problems that must be addressed when attempting to use uplink arrays for planetary radar applications, namely:

(1) Uplink array transmitters at Apollo operate near 7.2 GHz due to the uplink frequency allocation, whereas the GSSR transmitters have been allocated a frequency band at 8.5 GHz. This means that future uplink array transmitters need to cover both of these carrier frequencies.

(2) The operational downlink receivers operate at nominally 8.4 GHz, hence, their bandwidth must also be extended to cover the 8.5-GHz radar channel.

(3) Currently, the largest available 70-m aperture is used to transmit as well as receive Doppler-delay signals, providing the greatest possible gain for the radar receiver. For radar interferometry, two receivers separated by a large baseline are required instead of just one receiver for Doppler-delay imaging. Therefore, several receivers or receiver arrays separated by intercontinental distances may be required to process the weak planetary radar echoes; to reduce cost, these receivers should be designed for both the allocated radar carrier frequency and the communication frequency to provide the greatest flexibility.

Application of Uplink Arrays to Radio Science. There is considerable interest in radio science for uplink occultation experiments at Pluto and the outer solar system, where the increased EIRP provided by an uplink array would be very beneficial. The first uplink occultation experiment is planned for New Horizons at Pluto. It appears that if we can point the main beam to the spacecraft undergoing planetary occultation, then radio science requirements would be satisfied. Blind-pointing uplink array experiments designed to demonstrate the required array pointing accuracy would be useful to demonstrate array capabilities for future radio science applications.

VI. Summary and Conclusions

The first stable arraying of X-band carriers at interplanetary distances using up to three operational antennas has been demonstrated experimentally with the EPOXI spacecraft, on June 27, 2008. The experiment was carried out under realistic conditions at the Goldstone DSCC, using the 34-m BWG antennas located at DSS-24, DSS-25, and DSS-26 of the Apollo Complex. Doppler-compensated X-band carriers were transmitted to the spacecraft, and the concept of phase optimization to maximize received power was demonstrated through the use of a novel phase-ramping algorithm suitable for operational use with future geostationary satellites, or lunar and "nearby spacecraft" transponders. Power maximization was achieved after each phase ramp, validating the predicted array gain, and 50 No-Op commands were transmitted to EPOXI at the maximum rate, all of which were received and acknowledged by the spacecraft. Differential phase remained stable during the entire experiment, proving that the newly developed array frequency predicts were sufficiently accurate to enable future operational uplink arraying.

Acknowledgments

The authors would like to thank Chris Jacobs (Section 335) for providing accurate, VLBI-derived position vectors for the Apollo cluster phase centers that helped to generate the accurate array frequency predicts required for this experiment; Jonathan Walther contributed to this effort by developing the SPS predicts, and helped to analyze various experimental scenarios; Ryan Mukai (Section 337) deserves credit for significant contributions to the development of the LabView code used in the Phase Comparator and Control Assembly, and for simulating the Apollo cluster far-field array patterns of Figure 3. We thank Jim Taylor for invaluable help during the EPOXI planning process, and for providing the original WBAGC and CLA SNR data records from the telemetry delivery subsystem, shown in Figures 9 and 10.

References

[1] V. Vilnrotter and D. Lee, "Uplink Arraying Experiment with the Mars Global Surveyor Spacecraft," *The Interplanetary Network Progress Report*, vol. 42-166, Jet Propulsion Laboratory, Pasadena, California, pp. 1–14, August 15, 2006.
http://ipnpr.jpl.nasa.gov/progress_report/42-166/166F.pdf

[2] V. Vilnrotter, D. Lee, R. Mukai, T. Cornish, and P. Tsao, "Three-Antenna Doppler-Delay Imaging of the Crater Tycho for Uplink Array Calibration Applications," *The Interplanetary Network Progress Report*, vol. 42-169, Jet Propulsion Laboratory, Pasadena, California, pp. 1–17, May 15, 2007.
http://ipnpr.jpl.nasa.gov/progress_report/42-169/169D.pdf

[3] *DSMS Telecommunications Link Design Handbook*, DSN No. 810-005, Rev. E, Module 104: "34-m BWG Stations Telecommunications Interfaces," Rev. C, Jet Propulsion Laboratory, Pasadena, California, September 19, 2008.
http://eis.jpl.nasa.gov/deepspace/dsndocs/810-005/

[4] C. H. Acton, "Ancillary Data Services of NASA's Navigation and Ancillary Information Facility," *Planetary and Space Science*, vol. 44, no. 1, pp. 65–70, 1996.

[5] L. Paal, R. Mukai, V. Vilnrotter, T. Cornish, and D. Lee, "Ground System Phase Estimation Techniques for Uplink Array Applications," *The Interplanetary Network Progress Report*, vol. 42-167, Jet Propulsion Laboratory, Pasadena, California, pp. 1–15, November 15, 2006.
http://ipnpr.jpl.nasa.gov/progress_report/42-167/167H.pdf

[6] V. Vilnrotter, R. Mukai, and D. Lee, "Uplink Array Calibration via Far-Field Power Maximization," *The Interplanetary Network Progress Report*, vol. 42-164, Jet Propulsion Laboratory, Pasadena, California, pp. 1–16, February 15, 2006.
http://ipnpr.jpl.nasa.gov/progress_report/42-164/164D.pdf

[7] V. Vilnrotter, D. Lee, R. Mukai, T. Cornish, and P. Tsao, "Doppler-Delay Calibration of Uplink Arrays via Far-Field Moon-Bounce Power Maximization," *Proceedings of the 11th International Space Conference of Pacific Basin Societies (ISCOPS) Conference*, Beijing, China, May 15, 2007.